

What is claimed is:

1. An apparatus comprising:

a plurality of processing units;

a monitor to obtain a plurality of monitor values from said plurality of processing units, wherein said monitor is to transfer a process from a first processing unit of said plurality of processing units to a second processing unit of said plurality of processing units in response to said plurality of monitor values.

2. The apparatus of claim 1 wherein said monitor is to transfer the process from the first processing unit to the second processing unit in response to a first one of said plurality of monitor values being greater than a second one of said plurality of monitor values over a period of time.

3. The apparatus of claim 2 wherein said monitor obtains a monitor value by at least one of the set consisting of:

receiving a temperature indicator;

estimating an activity level;

receiving a power consumption estimate.

4. The apparatus of claim 1 wherein each of said plurality of processing units is one of a set consisting of:

10020568-120601

- 3 a core of a multi-core processor;
- 4 an execution unit of a processor;
- 5 a separate processor unit.

1 5. The apparatus of claim 1 wherein said monitor is further to increase and decrease a  
2 voltage level depending on a total power consumption or temperature level of said  
3 plurality of processing units.

1 6. The apparatus of claim 1 wherein said monitor comprises:  
2 an exchange module to exchange processes between ones of said plurality of  
3 processing units.

1 7. The apparatus of claim 1 wherein said monitor comprises:  
2 a move module to move one process from one of said plurality of processing units  
3 to another one of said plurality of processing units that is idle.

1 8. The apparatus of claim 6 wherein said monitor further comprises:  
2 a move module to move one process from one of said plurality of processing units  
3 to another one of said plurality of processing units that is idle;  
4 a sum module to throttle processing of one or more of said plurality of processing  
5 units if a sum total of power consumption of said plurality of processing units  
6 exceeds a selected total power consumption metric;  
7 a shutdown module to shut down one or more of said plurality of processing units

10020568.120501

8 in a low power mode.

1 9. The apparatus of claim 1 further comprising:

2 a cache coupled to said plurality of processing units, wherein said monitor is to  
3 swap processes between said first processing unit and said second processing  
4 unit by saving a first plurality of state variables from said first processing unit  
5 in said cache and saving a second plurality of state variables from said second  
6 processing unit in said cache and restoring said second plurality of state  
7 variables to said first processing unit from said cache and restoring said first  
8 plurality of state variables to said second processing unit from said cache.

1 10. The apparatus of claim 9 wherein said cache, said first processing unit, and said  
2 second processing unit are integrated on a single integrated circuit die, and wherein  
3 said cache is physically positioned between said first processing unit and said second  
4 processing unit.

1 11. The apparatus of claim 1 wherein said first processing unit and said second  
2 processing unit are coupled to receive power from different power wells and are  
3 capable of being independently operated at different voltages and frequencies under  
4 control of the monitor.

1 12. A multi-core processor comprising:

2 a first core having first instruction fetch and execute logic and a plurality of first

3 core state variables;  
4 a second core having second instruction fetch and execute logic and a plurality of  
5 second core state variables;  
6 a cache that is accessible to both said first core and said second core for read and  
7 write accesses;  
8 a monitor to monitor temperature and/or power consumption of said first core and  
9 said second core, and, in response to a selected metric being reached by one of  
10 said first core and said second core, to trigger storage of said plurality of first  
11 core state variables and said plurality of second core state variables in said  
12 cache and restoring of said plurality of second core state variables to said first  
13 core and restoring of said plurality of first core state variables to said second  
14 core.

1 13. The multi-core processor of claim 12 wherein said first core and said second core are  
2 coupled to independently controllable power supplies which are controllable by said  
3 monitor.

1 14. The multi-core processor of claim 13 wherein said first core and said second core are  
2 independently operable at different frequencies under control of said monitor.

1 15. The multi-core processor of claim 14 wherein said selected metric comprises at least  
2 one of the set consisting of:  
3 a level of processing activity;

4 a temperature level.

1 16. A system comprising:

2 a plurality of processing units, each processing unit to track its power

3 consumption, and to support a process move procedure;

4 a monitor to receive monitor information from each of said plurality of processing

5 units and to re-allocate processes to different ones of said plurality of

6 processing units in response to the monitor information received from the

7 plurality of processing units;

8 a memory coupled to said plurality of processing units to store instructions for

9 execution by said plurality of processing units.

1 17. The system of claim 16 wherein said monitor comprises:

2 a power-aware scheduler to schedule tasks for specific ones of said plurality of

3 processing units in response to said monitor information received from said

4 plurality of processing units.

1 18. The system of claim 17 wherein said power-aware scheduler is chosen from the set

2 consisting of:

3 an operating system scheduler that is stored in said memory during operation;

4 a hardware scheduler.

1 19. The system of claim 16 wherein said monitor comprises:

2 an exchange module to exchange processes between ones of said plurality of  
3 processing units;  
4 a move module to move one process from one of said plurality of processing units  
5 to another one of said plurality of processing units that is idle;  
6 a sum module to throttle processing of one or more of said plurality of processing  
7 units if a sum total of power consumption of said plurality of processing units  
8 exceeds a selected total power consumption metric;  
9 a shutdown module to shut down one or more of said plurality of processing units  
10 in a low power mode.

1 20. A method comprising:

2 monitoring power consumption of a plurality of processing units;  
3 swapping processes between said plurality of processing units in response to  
4 monitoring power consumption of said plurality of processing units.

1 21. The method of claim 20 wherein swapping comprises:

2 exchanging processor state data via a cache memory.

1 22. The method of claim 20 further comprising:

2 moving a process from a first one of said plurality of processing units to an idle  
3 one of said plurality of processing units in response to monitoring power  
4 consumption of said plurality of processing units.

1 23. The method of claim 21 further comprising:

2 reducing power consumption of one or more of said plurality of processing units

3 in response to a sum of power consumed exceeding a selected total power

4 consumption metric;

5 increasing power consumption of said plurality of processing units in response to

6 the sum of power consumed being less than a second selected total power

7 consumption metric.

1 24. The method of claim 23 further comprising:

2 periodically rearranging processes among said plurality of processing units.

1 25. The method of claim 20 further comprising:

2 independently controlling voltages and frequencies for said plurality of processing

3 units in response to monitoring power consumption of the plurality of

4 processing units.

1 26. An apparatus comprising:

2 a plurality of processing units;

3 a module to periodically transfer processes from a first processing unit from

4 said plurality of processing units to a second processing unit from said

5 plurality of processing units.

1 27. The apparatus of claim 26 further comprising a thermal monitor to independently  
2 control voltage levels of said plurality of processing units in response to a plurality of  
3 temperature levels of said plurality of processing units.

1 28. The apparatus of claim 27 wherein said thermal monitor is also to independently  
2 control clock frequencies for said plurality of processing units in response to said  
3 plurality of temperature levels.

1 29. An article comprising a machine readable medium storing a plurality of instructions  
2 which, if executed by a machine, cause the machine to perform operations  
3 comprising:

4 monitoring power consumption and/or thermal levels of a plurality of processing  
5 units;

6 swapping processes between said plurality of processing units in response to  
7 monitoring power consumption of said plurality of processing units.

1 30. The article of claim 29 wherein swapping comprises:  
2 exchanging processor state data via a cache memory.

1 31. The article of claim 29 wherein said operations further comprise:  
2 moving a process from a first one of said plurality of processing units to an idle  
3 one of said plurality of processing units in response to monitoring the power  
4 consumption of said plurality of processing units.



10020556-120601

1 32. The article of claim 30 wherein said operations further comprise:  
2 reducing power consumption of one or more of said plurality of processing units  
3 in response to a sum of power consumed exceeding a selected total power  
4 consumption metric;  
5 increasing power consumption of said plurality of processing units in response to  
6 the sum of power consumed being less than a second selected total power  
7 consumption metric.

1 33. The article of claim 32 wherein said operations further comprise:  
2 periodically rearranging processes among said plurality of processing units.

1 34. The article of claim 29 wherein said operations further comprise:  
2 independently controlling voltages and frequencies for said plurality of processing  
3 units.

1 35. A method comprising:  
2 monitoring temperature levels of a plurality of processing units;  
3 swapping processes between said plurality of processing units in response to  
4 monitoring temperature levels of said plurality of processing units.

1 36. The method of claim 35 wherein swapping comprises:  
2 exchanging processor state data via a cache memory.

- 1 37. The method of claim 35 further comprising:
- 2 periodically rearranging processes among said plurality of processing units.

10020568 130601